arXiv:cond-mat/0010104v1 [cond-mat.dis-nn] 6 Oct 2000

MAXENT QUERIES AND SEQUENTIAL SAMPLING

PETER RIEGLER
ABB Corporate Research, Speyerer Straße 4, D-69115 Heidelberg, Germany[‡]

AND

NESTOR CATICHA
Instituto de Física, Universidade de São Paulo, CP 66318 São Paulo, SP, CEP05389-970 Brazil[§]

Abstract. In this paper we pose the question: After gathering N data points, at what value of the control parameter should the next measurement be done? We propose an on-line algorithm which samples optimally by maximizing the gain in information on the parameters to be measured. We show analytically that the information gain is maximum for those potential measurements whose outcome is most unpredictable, i.e. for which the predictive distribution has maximum entropy. The resulting algorithm is applied to exponential analysis.

Key words: exponential analysis, design of experiment, Laplace transform, learning by queries

1. Introduction

Experiments usually measure the response of a physical system to an input. The measured set of outputs together with the corresponding inputs is then used to determine the parameters of some assumed model. For instance, for determining the electric resistance one would measure the resulting current for several values of the applied voltage. The voltage acts as a sort of control parameter. For time series the corresponding control parameter would be the time instances at which the signal is sampled.

The most common procedure for measuring physical parameters is to sample the underlying physical signal equidistantly with respect to the control parameter. However, it is often not clear how the accuracy of the parameters to be measured depends on the range over which the control parameter is varied. This is typical for signals of the form $\exp(-t/\tau)$, which decay exponentially in the control param-

[‡]Email: peter.riegler@de.abb.com
[§]Email: nestor@if.usp.br

eter t. For instance, if one happens to choose the maximum value of the control parameter to be much larger than the unknown decay constant τ many samples would obviously be "pure noise", thus reducing the accuracy of the estimate of τ.

The situation is particularly severe for sums of exponentials. There, the sampling rate clearly determines the smallest resolvable decay constant. Since for equidistant sampling those parts of the signal with larger decay constants are sampled more often, the estimates become the more accurate the larger the decay constant. Hence, a lot of sampling is wasted in order to estimate the easily to determine larger decay constants.

An experimentalist measuring a spectrum of decay constants consequently faces the question how to choose the values of the control parameters in order to obtain accurate measurements while restricting the maximum number of experiments to N. This task is basically one of design of experiment, a discipline well-known in statistics [1]. Due to the nonlinear dependence on the decay constants this is a nontrivial task, however. It would involve an optimization with respect to the control parameters of the N measurements to be done.

In this paper, we take a simplified approach by posing the question: After gathering N data points, at what value of the control parameter should the next measurement be done? We propose an on-line algorithm which samples optimally by maximizing the gain in information on the parameters to be measured. By this, the experiment is designed "sequentially" or "on-line" rather than completely from the onset.

Due to its design, the algorithm will choose at each step of the experiment the value of the control parameter at which the next experiment should be performed. In Section 2 we choose the entropy of the posterior as an information measure and introduce the algorithm for sequential MaxEnt sampling. We show analytically that the algorithm reduces to finding that value of the control parameter for which the width of the predictive distribution is maximum.

Section 3 gives an application of the algorithm to exponentially decaying signals and Section 4 compares the performance of the algorithm with equidistant sampling. In Section 5 we apply the algorithm to exponentially decaying signals having more than one decay constant. Finally, Section 6 summarizes the results and provides a conclusion.

2. MaxEnt formulation of sequential sampling

Assume we have determined a discrete data set $D_N = \{(t_1, s_1), ..., (t_N, s_N)\}$, sampled from $s(t)$ at discrete instances of the control parameter $t_1, ..., t_N$, with a model equation

$$s_i = s(t_i) = f(t_i) + \xi_i, \tag{1}$$

where $f(t)$ is the signal and ξ represents noise in the problem. Usually one has some parameterized form of $f(t)$ and the aim of the experiment is to determine (at least some of) these parameters. Having obtained D_N, the knowledge of these

parameters is contained in the posterior distribution

$$P(\boldsymbol{\beta}|D_N) = \frac{P(D_N|\boldsymbol{\beta})P(\boldsymbol{\beta})}{P(D_N)}, \tag{2}$$

where $\boldsymbol{\beta}$ denotes the whole set of parameters to be determined. For the prior probability of the noise $P(\xi) = P(D_N|\boldsymbol{\beta})$ we take the least informative prior [2], i.e. a Gaussian of zero mean and variance σ^2.

The aim of the experiment is to maximize the information about $\boldsymbol{\beta}$ during the course of the experiment. Using the negative entropy of the posterior as a measure of information this amounts to minimizing the entropy

$$S(D_N) = -\int d\boldsymbol{\beta} P(\boldsymbol{\beta}|D_N) \log P(\boldsymbol{\beta}|D_N). \tag{3}$$

The most informative experiment would then be defined by those values of the control parameter t which satisfy $S(D_n) = min$, implying

$$\frac{\partial S(D_N)}{\partial t_k} = 0 \quad \text{for all} \quad k = 1, ..., N. \tag{4}$$

In general this equation will be too difficult to solve. Moreover, the optimization (4) is not advisable, since it relies on a fixed model for $f(t)$. In general, there might be different optimal experiments for different models.

These difficulties can be avoided if one chooses to design the experiment in a sequential way. Instead of determining the values of all t_i one could ask the question: After gathering N data points, at what value of the control parameter t_{N+1} should the next measurement be done in order to gain as much information as possible about the parameters $\boldsymbol{\beta}$ to be estimated?

The answer to this is obtained by varying $S(D_{N+1}) = S\left(D_N \cup (t_{N+1}, s_{N+1})\right)$ with respect to the value of the control parameter t_{N+1} at the next experiment. That is, one has to determine the minimum of

$$\overline{S}(t_{N+1}) = \int ds_{N+1} S(D_{N+1}) P(s_{N+1}|D_N, t_{N+1}) \tag{5}$$

with respect to t_{N+1}. Note that one has to average out the unknown outcome s_{N+1} of the next experiment using the predictive distribution $P(s_{N+1}|D_N, t_{N+1})$. The predictive reflects the knowledge about s_{N+1} given the previously obtained data D_N and the control parameter of the next experiment.

Our sequential approach to design optimal experiments is somewhat similar to so-called query learning in neural networks. There one searches for such training examples that speed up the learning process considerably [3,4,5]. Based on this similarity we name the value of the control parameter t_{N+1} which minimizes (5) the *query* t_q.

Applying Bayes rule one can rewrite the r.h.s. of (5) in the form

$$\begin{aligned} \overline{S}(t_{N+1}) \quad = \quad & \int ds_{N+1} P(s_{N+1}|D_N, t_{N+1}) \log P(s_{N+1}|D_N, t_{N+1}) \\ & - \int d\boldsymbol{\beta} P(\boldsymbol{\beta}|D_N) \log P(\boldsymbol{\beta}|D_N) + \frac{1}{2}\left(1 + \log 2\pi\sigma^2\right). \end{aligned} \tag{6}$$

Note that the only dependence on t_{N+1} is via the entropy of the predictive distribution

$$S_{pred}(\tau) = -\int ds_{N+1} P(s_{N+1}|D_N, \tau) \log P(s_{N+1}|D_N, \tau). \tag{7}$$

Hence, in order to minimize the entropy of the posterior with a new measurement $(t_q, s(t_q))$ and N old measurements D_N, one has to *maximize* the entropy of the predictive distribution. The reason for this is that t_q is that value of the control parameter for which the prediction is least secure. Since the derived sampling procedure is based on a maximum entropy criterion, we call it sequential MaxEnt sampling.

Two notes are in place here, concerning the implementation of the above algorithm: According to (6), only $P(s_{N+1}|D_N, t_{N+1})$ is needed in order to determine t_q. By definition

$$P(s_{N+1}|D_N, t_{N+1}) = \frac{P(D_{N+1})}{P(D_N)} = \frac{\int d\boldsymbol{\beta} P(\boldsymbol{\beta}|D_{N+1})}{\int d\boldsymbol{\beta} P(\boldsymbol{\beta}|D_N)}. \tag{8}$$

For model functions of the form

$$f(t) = \sum_{j=1}^{m} A_j G_j(t, \{\lambda\}) \tag{9}$$

the integration with respect to the model parameters $\boldsymbol{\beta} = A_j, \lambda, \sigma$ can be simplified if one treats the parameters A_j as nuisance parameters, see [2] for details. In the same manner, the variance of the noise in (1) can be integrated out analytically using Jeffreys prior. Hence, the only parameters that actually need to be integrated over in (8) are the λ, which enter $f(t)$ in a nonlinear manner, in general.

For sufficiently large values of N the predictive distribution can be well approximated by a Gaussian of width σ_{pred}. The most probable value of $s(t_{N+1})$ will then be $(2\pi\sigma_{pred}^2)^{-1/2}$, which depends on t_{N+1} via σ_{pred}. Maximization of S_{pred} then amounts to searching for that value of t_{N+1} for which the most probable value of $s(t_{N+1})$ is minimum, see Figure 1.

3. Application to a monoexponentially decaying signal

As a first example we apply the proposed sequential MaxEnt sampling procedure to an exponentially decaying signal with only one decay constant:

$$s(t) = A \exp(-\lambda t) + \xi. \tag{10}$$

As before, ξ denotes additive noise of variance σ.

It is obvious that the sampling technique proposed in Section 2 cannot be applied from scratch, i.e. without any experimental data at hand. For estimating λ, for instance, one needs to have at least two measurements at different time instances since one needs at least two points to fit a straight line to $\log s$. Here we use three initial measurements, because two measurements lead to a posterior of width 0 which is difficult to handle numerically.

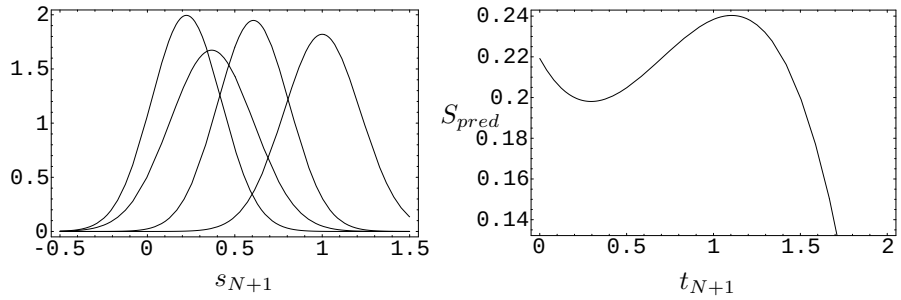

Figure 1. Left: Predictive distribution for various values of t_{N+1} (left to right: $t_{N+1} = 1.5$, 1.0,0.5,0). In the case shown the next measurement would be performed at $t_q = 1.1$. For this value the predictive distribution has the largest width and the corresponding entropy is maximum (right). The results have been obtained for a monoexponentially decaying signal as discussed in Section 3.

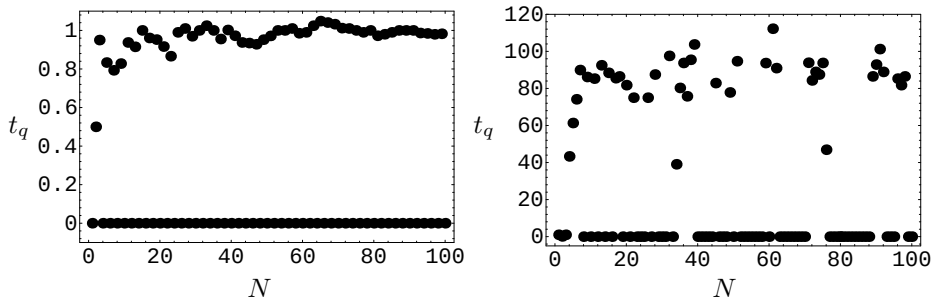

Figure 2. Simulation results for the sampling times t_q for $A = 1$, $\sigma = 0.1$. The decay constant has been chosen to be $\lambda = 1$ (left) and $\lambda = 0.01$ (right)

For the simulation results shown in the following, the required three initial measurements have been generated by drawing t_i, $i = 1..3$ from a uniform distribution in $[0, 2]$. In general, one would perform the initial measurements within an interval which is of the order of the decay time $\tau = 1/\lambda$. Note that this choice should be possible in principle, since an experimentalist will have some prior knowledge about the decay constant. However, such a choice is not mandatory, as is exemplified in Figure 2. There we show simulation results for $\lambda = 1$ and $\lambda = 0.01$. In the latter case the initial measurements are at times much smaller than the decay constant $\tau = 100$, yet the sequential MaxEnt sampling algorithm is capable of tracking the signal.

Figure 2 shows the evolution of sample times t_q for increasing length N of the experiment. We find the remarking behavior that the sampling times "oscillate" between $t_q = 0$ and $t_q \simeq 1/\lambda$. Note that the period is not strictly equal to one, since sometimes two or more consecutive measurements are at either $t_q = 0$ or

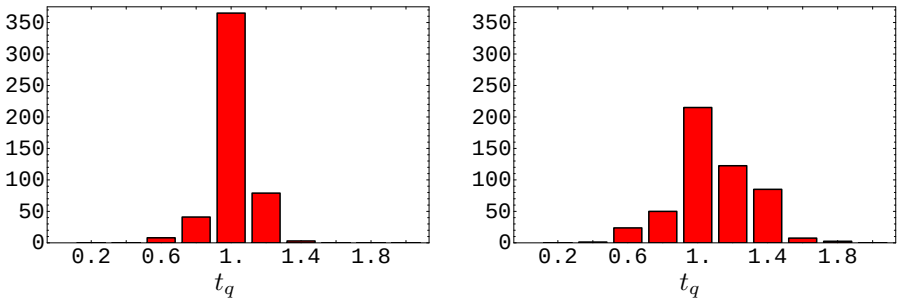

Figure 3. Distribution of sampling times $t_q \neq 0$ for an exponentially decaying signal of form (10) with $A = 1$ and $\lambda = 1$. The histograms have been obtained for 10 experiments of length $N = 100$. Left: $\sigma = 0.1$, right: $\sigma = 0.5$.

$t_q \simeq 1/\lambda$. The oscillating behavior can be traced back to the functional form of $S_{pred}(t_{N+1})$, cf. Figure 1. The entropy $S_{pred}(t_{N+1})$ has two maxima for $t \geq 0$, the absolute maximum determining t_q. By construction of MaxEnt sampling the next measurement will be placed at t_{global} and $S(t_{global})$ will decrease relative to $S(t_{local})$. After one or more measurements the predictive entropy at the hitherto global maximum will be lower than at the hitherto local maximum and the two maxima will change their roles.

From Figure 2 it is evident that the larger sampling times are not exactly at $t_q = 1/\lambda$. Instead, the values of t_q have some spread which increases with the variance σ^2 of the additive noise, see Figure 3.

In the course of the experiment the standard errors of the parameters to be estimated decrease with $1/\sqrt{N}$, asymptotically, as expected, cf. [2]. See Figure 4 for an example.

A note is in place here concerning the oscillating behavior of the control parameter. For time series, the control parameter t would be time. At first glance, it

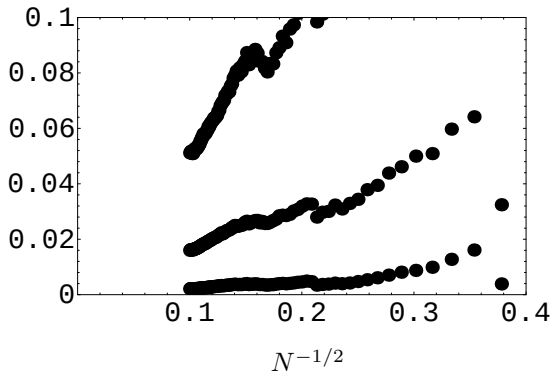

Figure 4. Standard errors of the estimates of λ, A, and σ (top to bottom) obtained via sequential MaxEnt sampling. Shown are the results for the same simulation as in Figure 2, left.

might seem impossible to perform sequential MaxEnt sampling since time can only be increased but not decreased. For many types of experiments, however, one can actually decrease the time of the next measurement by performing a new experiment. For instance, in measuring NMR relaxation times one would first magnetize the probe which defines the onset of the experiment, $t = 0$, and then measure the magnetization for several $t > 0$. For MaxEnt sampling, one would have to magnetize, measure at t_q, magnetize again, and so on. Hence, t_q can increase or decrease in consecutive experiments.

4. Comparison with equidistant sampling

It is instructive to compare the results obtained for sequential MaxEnt sampling with other sampling methods. The most common alternative method is probably equidistant sampling where the experiments are performed at N equidistantly spaced values of the control parameter. For a signal of form (10) the lowest of these values is obviously $t = 0$ while the largest t_{max} is up to the choice of the experimentalist.

Clearly, there is an optimal choice of t_{max}. For, choosing $t_{max} \gg 1/\lambda$ would result in sampling of noise only, while $t_{max} \ll 1/\lambda$ would give a poor estimate of λ as the noiseless signal $A\exp(-\lambda t)$ hardly varies in this range. This behavior is confirmed by simulations. In Figure 5 we show the standard errors of the Bayesian estimates of A, λ, σ for various values of t_{max}. As can be seen, the standard error of A increases with t_{max} while the standard error of λ has some minimum around $t_{max} \simeq 2/\lambda$; the standard error of the estimate of σ is basically independent of t_{max}.

Consequently, in choosing t_{max} there is a tradeoff between estimating A (leading to $t_{max} \to 0$) and λ (leading to $t_{max} \simeq 2/\lambda$). Besides this tradeoff, not knowing λ one could determine t_{max} only based on the prior knowledge on λ. In contrast to that, sequential MaxEnt sampling, as described in the previous section, does not require the setting of any maximum control parameter t_{max} and places its measurements automatically.

In addition, MaxEnt sampling leads to a much better resolution of the parameters to be estimated. This is exemplified in Figure 6, which compares the obtained standard errors of A, λ for MaxEnt and equidistant sampling.

5. Application to sums of exponentials

Exponentially decaying signals can be found in many physical phenomena as first order differential equations are among the most common in physics. However, the spectrum $\{\lambda_n, A_n\}$ of a multi-exponential signal

$$f(t) = \sum_n A_n \exp(-\lambda_n t) \tag{11}$$

is difficult to estimate. The reason is that the required inverse (discrete) Laplace transform is known to be an ill-posed problem [6,7]. A small amount of noise added

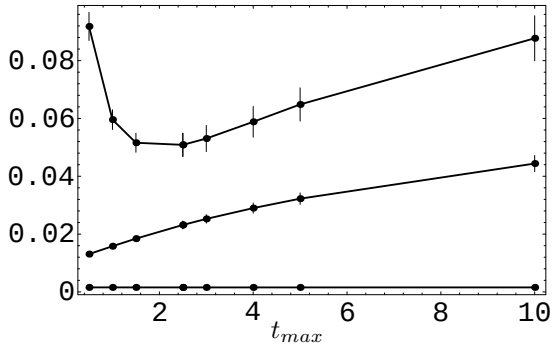

Figure 5. Standard errors of the Bayesian estimates of A, λ, σ of a signal of form (10). Results have been obtained for equidistant sampling of $N = 100$ points between $t = 0$ and t_{max}, averaged over 20 independent runs (same simulation settings as in Figure 2, left). Top to bottom: standard errors of λ, A, σ. Note that depending on the parameters to be estimated there is a different optimal choice for t_{max}. The standard error of σ is basically independent of t_{max}.

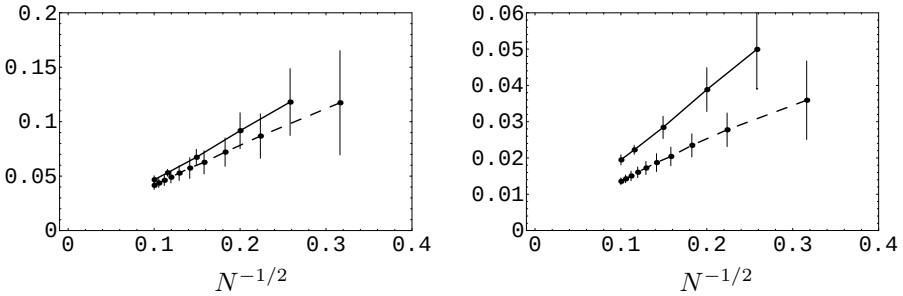

Figure 6. Comparison between sequential MaxEnt sampling (dashed line) and equidistant sampling (solid line). Shown are the standard errors for λ (left) and A (right). Simulation parameters are for the same settings as in Figure 5, averaged over 20 independent runs. Equidistant sampling has been performed with $t_{max} = 2.0$.

to $f(t)$ leads to a nonvanishing contribution to the spectrum even in the limit of a very small perturbation.

Moreover, the quality of the spectrum to be estimated from measurements depends strongly on the values of the control parameter t at which the measurements are performed [8]. For instance, choosing $t > 1/\lambda^*$ certainly rules out the detection of any $\lambda_n > \lambda^*$ in (11). Some good prior knowledge about the values of λ_n to be expected in the signal would be very helpful. But even such knowledge would leave open the question how to choose the sampling times t_i. For instance, in porous materials [9] the magnetization decays approximately like

$$f(t) = \sum_{n=1}^{\infty} \frac{A_0}{n^2} \exp\left(-\lambda_0 n^2 t\right), \tag{12}$$

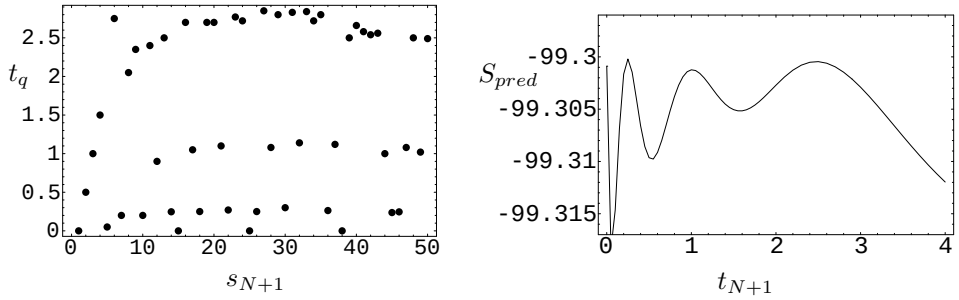

Figure 7. Left: Simulation results for the sampling times t_q for a signal of type (11) with two exponentials. Parameters are $A_1 = 1$, $A_2 = 1/4$, $\lambda_1 = 1$, $\lambda_2 = 4$, $\sigma = 0.01$. Note that the first four measurements are chosen equidistantly in order to start up the experiment. Right: Entropy of the corresponding predictive distribution after $N = 50$ measurements.

i.e. relaxation times and corresponding amplitudes rapidly decrease like $1/n^2$. How does one have to choose the sampling times in order to resolve the spectrum up to a desired n_{max}?

By construction of the algorithm, sequential MaxEnt sampling provides the answer to this question. The choice of n_{max} fixes the model function (11). Given some initial data, the best query t_q can be computed as described in Section 2. Figure 7 shows the simulation result for the values of t_q for a signal of form (12) with $n_{max} = 2$.

As in the monoexponential case, the values of the sampled control parameters oscillate between $t = 0$ and the characteristic time constants in the signal, i.e. $1/\lambda_1$ and $1/\lambda_2$ in our case. Note that in Figure 7 the signal is also sampled at times $t > 1/\lambda_1$ larger than the largest characteristic time constant The reason for this is probably due to the fact that we have not imposed any priors on the parameters λ. In order to rule out the presence of any $\lambda < 0$ one would clearly have to sample at large values of t.

Figure 7 also depicts the characteristic dependence of the predictive entropy on t_{N+1}. As in Section 3 the location of its maxima characterizes the sampled values of the control parameter t.

In Figure 8 we show the resulting posterior on λ (with A_i and σ treated as nuisance parameters) in order to compare with equidistant sampling. As can be seen, sequential MaxEnt sampling leads to lower error bars in the estimates of λ. In particular the width of the posterior in the large decay constant λ_2 is considerably larger for equidistant sampling. For signals of the form (12) this is exactly what one would expect since the term with the smaller decay constant λ_1 is sampled more frequently than the term decaying with λ_2. For sequential MaxEnt sampling, however, the latter term is sampled more frequently, cf. Figure 7.

Despite the good performance of sequential MaxEnt sampling it has to be stressed that it becomes impracticable in its current form for $n_{max} > 2$. For a low number N of experiments it will be difficult to find evidence for the decay constant

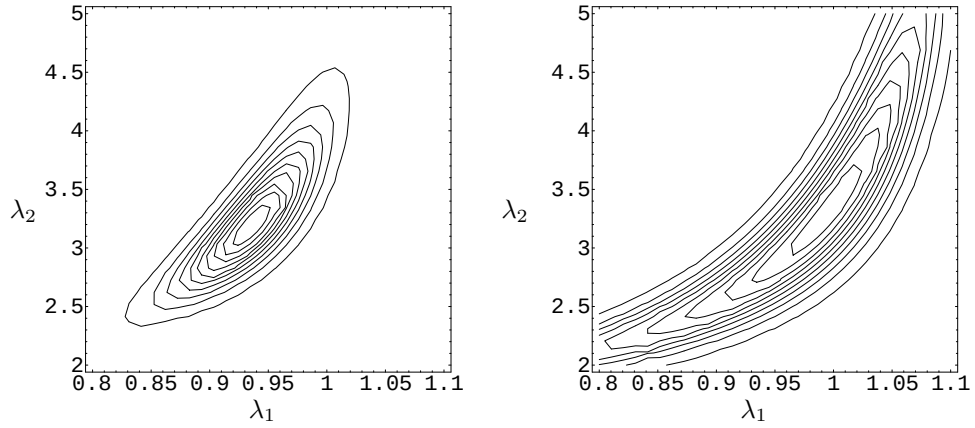

Figure 8. Contour plots of the posterior $P(\lambda_1, \lambda_2|D_{50})$. Left: posterior corresponding to the MaxEnt sampling experiment of Figure 7. Right: posterior as obtained after equidistantly sampling 50 data points between $t = 0$ and $t_{max} = 1.5$. The chosen value of t_{max} represents the optimal if one aims at highly accurate estimates of the parameters λ. In contrast to monoexponential sampling, however, the optimal choice of t_{max} hardly depends on the values of λ for intermediate values of t_{max}

$\lambda_{n_{max}}$ in (12). A low number of experiments N will result in a degeneracy in the estimates of some of the λ_i leading to divergences in the posterior. The remedy here would be to incorporate model selection as described in [2] into the algorithm. This would probably result in an algorithm which starts with the hypothesis $n_{max} = 1$ for low N and increases n_{max} as soon as the relative probability for such a hypothesis is larger.

6. Summary

Starting from the question at what value of the control parameter one should perform the next experiment after having performed N measurements, we have analytically derived a sequential sampling procedure which maximizes the information on the parameters to be estimated. Looking for the query t_{N+1} that leads to a maximum information gain (minimization of the posterior entropy) led us to the conclusion that one should find t_q that *maximizes* the entropy of the predictive distribution. The reason for this is that t_q is the value of the control parameter for which the prediction is least secure.

We have applied the constructed MaxEnt sampling procedure to the fitting of exponential signals. We found that the resulting queries oscillate between the time constants present in the signal. The width of the distribution of queries around the time constants depends on the variance of noise by which the signal is corrupted.

For purely monoexponential signals we have compared sequential MaxEnt sam-

pling with equidistant sampling and found a considerably better accuracy of the estimates of MaxEnt sampling. The application of sequential MaxEnt sampling to multiexponential signals is of particular interest because of their frequent appearance in science and the difficulties associated with the inverse Laplace transform. We have motivated that proper sampling helps to perform the inversion task. We have exemplified that sequential MaxEnt sampling is capable of tracking the spectrum quickly.

However, so far we have applied the algorithm only to an assumed constant number of exponentially decaying terms. This is a bit unrealistic, since typically there would be infinitely many such terms, like the NMR signals mentioned in Section 5. Moreover, the assumption of a large fixed number of exponential terms in the signal leads to degeneracies and hence divergences in the posterior. We hope that these divergences can be removed by hypothesizing a small number n_{max} of exponential terms at the onset of the experiment and increasing this number whenever the relative probability favors the hypothesis of a larger value of n_{max}.

Finally we note that due to the nonlinear dependence of the model functions on the decay constants λ results could only be obtained numerically. For models completely linear in the model parameters it is possible to do the calculation of query times t_q analytically [10], however. For such a case, the result is rather trivial, leading to $t_q \to \infty$.

The authors are grateful to S. Rabbani and C. Mendonça for stimulating discussions. P.R. acknowledges the hospitality extended to him at Universidade de São Paulo where this work has been initiated.

References

1. Tagouchi, *System of experimental design*, Quality Resources, White Plains, NY, 1991.
2. G. Bretthorst, *Bayesian Spectrum Analysis and Parameter Estimation*, Springer, Berlin, 1988.
3. W. Kinzel and P. Ruján, "Improving a network generalization ability by selecting examples," *Europhys. Lett.*, **13**, p. 473, 1990.
4. P. Sollich, *Asking Intelligent Questions – The Statistical Mechanics of Query Learning*, Ph.D. thesis, University of Edinburgh, 1995.
5. K. Fukumizu, "Statistical active learning in multilayer perceptrons," *IEEE Transactions on Neural Networks*, **11**, pp. 17–26, 2000.
6. B. Hüpper and E. Pollak, "Numerical inversion of the Laplace transform," *J. Chem. Phys.*, **110**, pp. 11176–11186, 1999.
7. V. N. Vapnik, *Statistical Learning Theory*, Wiley, New York, 1998.
8. A. A. Istratov and O. F. Vyvenko, "Exponential analysis in physical phenomena," *Rev. Sci. Instrum.*, **70**, pp. 1233–1254, 1999.
9. S. R. Rabbani and D. T. Edmonds, "Nuclear spin-lattice relaxation-time reduction in small particles," *Phys. Rev.* **B**, **50**, pp. 6184–6188, 1994.
10. unpublished results